

APPLICATION FOR PATENT

Inventors: YARON KASHAI and MATTHEW JOHN MORLEY

Title: SYNTHESIS OF VERIFICATION LANGUAGES

5 FIELD OF THE INVENTION

The present invention relates to the synthesis of verification specific languages, and in particular, to a method for compiling a program written in such a language into a target language, of which a preferred example is a HDL (hardware description language).

10

BACKGROUND OF THE INVENTION

Design verification is the process of determining whether an integrated circuit, board, or system-level architecture, exactly implements the requirements defined by the specification of the architecture for that device. Design verification for a device under testing (DUT) may be performed on the actual device, or on a simulation model of the device. For the purposes of explanation only and without intending to be limiting in any way, the following discussion centers upon testing which is performed on simulation models of the device.

As designs for different types of devices and device architectures become more complex, the likelihood of design errors increases. However, design verification also becomes more difficult and time consuming, as the simulation models of the design of the device also become more complex to prepare and to test.

The problem of design verification is compounded by the lack of widely generalizable tools which are useful for the verification and testing of a wide variety of devices and device architectures. Typical background art verification methods have often been restricted to a particular device having a specific design, such that the steps of preparing and implementing such verification methods for the simulation model must be performed for each new device.

The process of verifying a design through a simulation model of the device is aided by the availability of hardware description languages, or HDL, such as Verilog and VHDL. These languages are designed to describe hardware at higher levels of abstraction than gates or transistors. The resultant simulated model of the device can receive input stimuli in the form of test vectors, which are a string of binary digits applied to the input of a circuit. The simulated model then produces results, which are checked against the expected results for the particular design of the device. However, these languages are typically not designed for actual verification. Therefore, the verification engineer must write additional programming code in order to interface with the models described by these hardware description languages in order to perform design verification of the device. In particular, the *e* verification specific language is highly useful for design verification.

One example of a useful verification language is the *e* language which is provided by the Specman™ functional programming environment (Verisity Design, Inc., Mountain View, California, USA) and disclosed in U.S. Patent Application No. 6,182,258, filed on February 6, 1998, which is hereby

incorporated by reference as if fully set forth herein. Although the *e* language is highly useful for creating programs for design verification, it still has certain drawbacks. For example, like most software languages, *e* code does not have a clear static structure, and the dynamic functionality of the model is open to
5 interpretation at runtime, most notably with regard to timing. By contrast, hardware description languages have a more clear static structure and are less ambiguous with regard to issues such as timing. However, this latter type of language does not inherently provide the dynamic aspects of verification languages such as the *e* language.

10

SUMMARY OF THE INVENTION

None of the background art references teaches or suggests a method for automatically translating the *e* language into a hardware description language.

The present invention overcomes these drawbacks of the background art
15 by providing a method for synthesizing the *e* language, and thereby enabling the language to be compiled into a target language. This method enables the underlying control structure of the language to be determined, and then used to map the dynamic behavior of the language onto the target language as part of a static framework. The process of synthesizing any type of language causes at
20 least a portion of the implicit control structure of the software program to be constructed into the compiled output code, such that an additional scheduler or other type of runtime system may not be required. Therefore, the compiled output code should have a greater execution speed and should be operated more

efficiently than the software programs which are written in the source language itself.

According to the present invention, there is provided a method for automatically translating code written in a first language to a second target
5 language, the first language featuring dynamic behavior, wherein the second target language features a hierarchy of objects, such that relationships between the objects are determined by constraints in the first language, the method comprising: detecting an underlying control structure for the code in the first language; creating a static framework of resources for supporting the dynamic
10 behavior of the code written in the first language; and mapping the dynamic behavior to the second target language according to the static framework of resources and according to the underlying control structure.

The method of the present invention could also be described as a plurality of instructions being performed by a data processor, such that the method of the
15 present invention could be implemented as hardware, software, firmware or a combination thereof. For the present invention, a software application could be written in substantially any suitable programming language, which could easily be selected by one of ordinary skill in the art. The programming language chosen should be compatible with the computing platform according to which the
20 software application is executed. Examples of suitable programming languages include, but are not limited to, C, C++ and Java.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention is herein described, by way of example only, with reference to the accompanying drawings, wherein:

FIG. 1 is a flowchart of an illustrative method according to the present invention;

FIG. 2 shows an example of “pseudocode” for performing a breadth-first elaboration according to the present invention;

FIG. 3 shows an example of the elaboration process with regard to implementation with the *e* language according to the present invention;

FIG. 4 shows an exemplary graph according to the present invention which results from the execution of the code of Figure 2;

FIG. 5 shows the exemplary graph of Figure 4 after being unfolded;

FIG. 6 shows an example of *e* language code input for creating a sequential control flow graph according to the present invention;

FIG. 7 shows the resultant sequential control flow graph from Figure 6 according to the present invention;

FIG. 8 shows an exemplary unrolled graph after the process of unrolling is performed on the sequential control flow graph of Figure 7 according to the present invention;

FIG. 9 shows an example of the segmentation of the client process for the sequential control flow graph according to the present invention;

FIG. 10 shows exemplary pseudocode for an exemplary annotation process with regard to tagging according to the present invention;

FIG. 11 shows a simple two process example as written in the *e* language;

FIG. 12 shows a resultant elaboration graph, created according to the method of the present invention;

FIG. 13 shows an exemplary segmented scheduling constraint graph;

5 FIG. 14 shows a derived segmented scheduling constraint graph

DESCRIPTION OF THE PREFERRED EMBODIMENTS

6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2

be mapped onto the target language, such that the execution of the compiled code accurately represents the dynamic behavior of the verification language itself.

Although the *e* temporal language was not designed for synthesis, it does have a general structure which can be related to the behavior of hardware devices, to a much greater extent than “pure” software languages such as Java and C++. Therefore, the *e* language is the preferred verification language for the present invention, although as previously noted, it is not the only such verification language which is useful for the synthesis process.

Referring now to the drawings, Figure 1 is a flowchart of an exemplary method for synthesis according to the present invention. As shown, the method first involves “elaborating” the verification language, in a process which is termed herein “elaboration”. The process of elaboration involves the allocation of sufficient state holding elements to represent all possible behavior which may be invoked by the *e* language program. In particular, structures which are present prior to execution are considered to be static, while objects allocated “on the fly” during execution are considered to be transient. The latter objects may be handled at run time by data transfer, as all required structures should be represented in the compiled language prior to execution.

In order to represent these structures, preferably the static structure of the verification language is expanded according to predefined semantic constructions, thereby resulting in compiled code of a predictable construction. Although alternatively the dynamic properties of the verification language could be captured by executing the verification language, such that the dynamic aspects

of the language are literally operated, such an alternative approach has a number of disadvantages. For example, higher level information regarding the roles of scalars and objects is lost in the alternative approach, such that further analysis of the dynamic structure is not possible. Thus, for the preferred embodiment of the
5 present invention, the static structure is determined according to predefined semantic constructions.

These predefined semantic constructions optionally and more preferably include nesting of types and the determination of structural constraints. The process of construction according to nesting of types means that elaboration
10 always starts from a top unit, and then recursively includes descendants of that unit. Structural constraints are embedded within the type declarations, and are used to impose relationships between parts of the structure, as well as for imposing initial values and constant values on fields. Timing-dependent, or *when inheritance* constructions further complicate this process, since these
15 constructions may require the determination of an entire subtree of possibilities if the exact construction cannot otherwise be determined.

As shown with regard to Figure 1, the method of the present invention preferably starts with parsing the input program written in the source language and creating an abstract syntax tree, as known in the art. The process continues
20 with the creation of an elaboration graph, which is the start of the elaboration process. . For the preferred graph embodiment, each instance of an object or a scalar is represented by a node. All graph nodes contain their type information. Preferably, there are three major types of nodes: scalar nodes, which may contain

initial or constant values; struct nodes which reference nested instances; and list nodes, which reference list items, list size and index roles. All nodes have a parent reference which is decided according to the creation order (as described in greater detail below). Directed edges mark references. Edges are tagged with the
5 name which is used to access the target when positioned at the source.

The elaboration process is preferably performed with a breadth-first traversal, as shown with regard to Figure 1. An example of "pseudocode" for performing such a breadth-first traversal is shown with regard to Figure 2. The breadth-first traversal involves the collection of structural constraints, and the
10 application of such constraints at each level of the traversal process. Constraints preferably assign values and unify instances. For the latter activity, more preferably only the instance closer to the root is retained, in order to avoid infinite recursion.

If a constraint cannot be applied at that level, for example because the
15 constraint contains an argument which has not yet been invoked, then preferably it is collected into an "unsatisfied constraint list". Also preferably, the breadth-first search includes the process of repeatedly examining this list, in an attempt to apply one or more of the constraints at each level of the search.

Next, a unique parent is identified for each instance, which is the first or
20 topmost entity that referenced the instance. It should be noted that if multiple references are made from the same level of the search, the parent relationship may optionally be considered to be arbitrary.

According to preferred embodiments of the present invention, special consideration is given to maintaining any type of symmetry which is inherent in the verification language, at least to the extent possible while also synthesizing the program code. The preferred implementation of the present invention with
5 the *e* language is therefore helpful for the maintenance of symmetry, since certain features of the *e* language are helpful for detecting and maintaining symmetry. For example, lists are widely used in the *e* language, often in a manner which preserves symmetry. Lists themselves are inherently symmetric types, in which a representative item can stand for all items, thus simplifying the analysis. The
10 method of the present invention further maintains such symmetry by more preferably providing list nodes which have dual representation, with a symmetric list item representative which stands for all items, and a list of all outstanding items which are distinguished from the symmetric list item. According to the present invention, more preferably all lists are initially assumed to be symmetric,
15 with the process of synthesis switching from the symmetric representation to the explicit representation under certain conditions which break symmetry, as described in greater detail below.

Another example of maintaining symmetry is implemented with regard to scalars. For example, a scalar which is used as an index into a symmetric array.
20 Such a scalar does not break symmetry while being used as an index, as long as the scalar is not assigned a particular value.

These processes are preferably repeated in order to create the graph, thereby forming the elaboration process. Exemplary pseudocode for such a

process is shown with regard to Figure 2. As shown, *L* is the set of nodes in the current search level, while *C* is the unsatisfied constraint set. The function *apply()* takes the set of unsatisfied constraints, attempts to modify the structure to match the constraints and returns the set of constraints which remained
5 unsatisfied. The function *typeof()* receives an instance reference and returns a type descriptor, while the function *refto()* receives a type descriptor and returns a graph reference edge. The function *makenode()* receives a reference and returns the graph node corresponding to the instance.

As shown in Figure 2, the exemplary pseudocode algorithm starts with a
10 reference to the top unit. For each reference *r* in the set of nodes *L*, a type description is first determined. Next, the node corresponding to the reference is returned by the function *makenode()*. Constraints for that type description are added to the list of unsolved constraints *C*. Each field in the type description is then examined with regard to the list of unsolved constraints, in order to
15 determine whether any of the constraints may be applied. In addition, a new reference edge is also returned, for connecting the new graph node to the remainder of the graph. This process is then preferably repeated, more preferably until all of the constraints have either been applied or are considered to be unsolvable.

20 Figure 3 shows an example of the elaboration process with regard to implementation with the *e* language. This exemplary process includes nested structs and lists of structs, although methods have been omitted for brevity. An exemplary graph according to the present invention which results from the

execution of this code is shown in Figure 4. Boxes represent *e* structs, ellipses represent scalars and octagons represent lists. Dashed links represent special roles, such as indices.

The exemplary program contains a top level object called *sys*, which
5 contains a nested list of objects, called *cl_list*, the objects being of type *cl*. The constraints within *sys* determine *cl_list* to contain 4 elements, with each element's *id* field being set to the element's index in the list. *Sys* also contains an object called *arb* of type *arb*. The *cls* field of this object is constrained to reference the *cl_list* list mentioned above. The *arb* type is shown to contain a list of *cl type*
10 called *cls* and an integer field called *data*. The type *cl* is shown to contain multiple scalar fields and an *arb* field of type *arb*. The *arb* field is constrained to be the same as the *arb* field under *sys*.

Figure 4 shows the static structure that is extracted by the elaboration process. Using the constraints under *sys* and under *cl*, the process resulted in a
15 single list of *cl* type elements and a single *arb* object in spite of what seems to be an infinitely recursive structure.

Turning again to Figure 1, the process of elaboration is preferably completed when all of the constraints have either been applied or are considered to be unsolvable. However, optionally and more preferably, as shown, the
20 completed graph is unfolded in order to more completely represent the verification language structures. For example, the completed graph is more preferably unfolded when the symmetrical properties of lists are broken. An example of a situation in which the symmetry of a list is broken occurs when the

objects in a list have embedded Time Consuming Methods (which are independent threads of control or processes, termed “TCMs” herein) or temporal expressions. Since these processes have an internal state, and since these processes are likely to be at different states per instance, the list is no longer
5 symmetric.

Figure 5 shows the exemplary graph of Figure 4 after being unfolded, which in this example is required since the *cl* struct has a TCM in it. The unfolded graph includes all possibilities for this struct, since the state of the struct cannot be predetermined during the elaboration process.

10 Turning again to Figure 1, the exemplary method shows the analysis of the control structure, or control flow analysis. Although this analysis is shown as a separate and sequential stage which follows the elaboration stage, in fact the analysis of the control structure may optionally be performed in parallel with the process of elaboration. The analysis of the control structure is preferably
15 performed with regard to the methods and temporal expressions which are part of the *e* type declarations. In addition, the various instantiations of TCMs and events are preferably included in the analysis.

The end result of control flow analysis is to determine for each data processing action which conditions and triggers may cause that action to execute.
20 For the preferred implementation of the present invention for compilation into a hardware design language as the target language, the semantics may be assumed to be synchronous and cycle based, such that boolean conditions and temporal triggers may be considered to be similar. However, preferably temporal triggers

have computational priority over boolean expressions, for the preferred implementation of the present invention to correctly replicate the semantics of the *e* language. The only synchronous trigger remains the system clock, and all non-clock, temporal triggers and boolean conditions are termed “guards” herein.

- 5 Therefore, control flow analysis should enable a list of guards to be determined for governing the execution of each action, thereby resulting in the creation of a sequential control flow graph.

Turning again to Figure 1, control flow analysis is preferably initiated by first computing the triggering structure for each process. This computation is
10 optionally and preferably performed by capturing the flow of control from the *struct* level to the TCM, or the temporal expression. The flow of control is captured by scanning the *run()* method for any *start* actions, handling spawning inside of TCMs, handling *on* methods (which are procedures executed when particular temporal expressions succeed), and so forth.

- 15 Next, the guards are computed for each action of each process, by analyzing the action sequence for each method, in order to determine non-clock, temporal triggers and boolean conditions. Optionally and more preferably, method calls cause the control flow of the called method to be instantiated in place of the call itself.

- 20 Once the guards have been determined, preferably the sequential control flow graph is created. This graph groups actions into basic blocks. Basic blocks may optionally be nested within guards, while guards can optionally be nested as well as chained. Actions which occur over a period of time, or are otherwise

time-consuming such as *wait* and *sync*, are preferably represented by a special *wait node* which breaks basic blocks (as previously described, these blocks may optionally be segmented by actions which block processes, such as *wait*).

The sequential control flow graph is more preferably created from at least
5 three types of nodes: basic nodes, guard nodes and wait nodes. Basic nodes represent a list of actions, none of which is time-consuming or is a control flow action. Guard nodes represent a branch point, and have a boolean expression which is a guard, a positive branch which is taken if the expression evaluates to true and a negative branch which executes otherwise. A label on the guard can
10 mark it as a loop or an *if then else* action. A wait node is similar to a basic node, except that it contains a single temporal expression or other time consuming action, such as communicating through a port (and no other actions).

According to preferred embodiments of the present invention, the nodes of this graph retain references to the abstract syntax tree nodes, so that code
15 generation can extract additional information. The entire graph is associated with the particular type of node of the abstract syntax tree.

An example of *e* language code input for creating the sequential control flow graph is shown in Figure 6.

The code example contains a list of clients of type *cl* nested under the top
20 object called *sys*. An arbiter object of type *arb* is nested under *sys* as well. The clients attempt to transfer data to a shared data field nested within the arbiter. While the clients operate concurrently, the access to the common field must be performed sequentially to prevent data corruption. Sequential access is preferably

achieved by a protocol that executes between the arbiter and each client. When the client's data is ready for transmission, signaled by *drdy* being set by an external agent (not included in the sample code), the client asserts *xreq*. It then waits for the arbiter to notice the request and acknowledge by setting *xgrt*. At this point the data is transferred and the client deasserts request, which will cause the arbiter to deassert *xgrt*. Once the transfer is completed, the client deasserts *drdy*, to signal the termination for the external agent.

The resultant sequential control flow graph is shown in Figure 7. Boxes represent basic blocks, diamonds represent guards and octagons represent time consuming actions.

As the sequential control flow graph is created, context is more preferably established for at least one action. Certain actions, such as iterators and method calls, create local context. This local context is needed in order to interpret expressions and to locate references on the elaboration graph. The context for each action is preferably established by passing a list of name-reference pairs to the control flow extraction process. Additions are preferably made to the list each time that an action creates local variables. The list is then "popped", or non-relevant variables removed, each time that the process moves out of the scope of these variables. Optionally and more preferably, the context is saved on each wait node. Since execution may only commence at these nodes, the context information is more efficiently retrieved from these nodes.

According to optional but preferred embodiments of the present invention, the control flow extraction process may uncover at least one malformed control

structure, which the user then preferably alters manually. One example of such a malformed control structure is a cycle which contains only conditional *wait* actions. Since there is an infinite number of ways to unwind the cycle, such code cannot be represented as a finite list of segments, which is a fundamental
5 requirement of the code generation phase. Thus, the user is informed of the existence of such structures, and is preferably required to modify the code such that the loop is either wait free, or contains at least one unconditional wait.

Loops are then preferably unrolled, where needed. Some loops may need to be unrolled either because the loop structure is malformed as previously
10 described, or due to unfolding of structures to which the loop refers. In cases of malformed control structures, the user preferably manually alters the code before the loop is unrolled. With regard to the unfolding of control structures to which the loop refers, one example of a situation in which such unfolding occurs is for unfolding of lists due to breaking of symmetry. Certain processes may use
15 iterators to access the list; however, once the list is unfolded, iterative access is typically no longer supported, since the list structure has been altered. Therefore, the iterator itself must be unfolded in order for the list to be properly accessed.

The process of unrolling requires some support in terms of context value assignment. For example, a *for each* action adds variables such as *it* and *index* to
20 the context. When unrolled, these variables become constants whose scope is a single iteration. This process is supported by a fourth sequential control flow graph node, which is called a context node. The context node adds a scope with

either variables or constants to the context. A context node is added at the top of each iteration of the unrolled structure.

An example of an unrolled loop is described with regard to the example of Figure 6. In that example, the TCM nested in the client objects *cl* causes the list to be unfolded. Since the arbiter process *switch()* accesses the list objects, the list iterator must be unrolled. The process of unrolling is performed on the sequential control flow graph, with the resultant new graph for the unrolled arbiter process being shown in Figure 8. The dashed boxes in Figure 8 represent context nodes.

Next, the segmentation of each process is more preferably computed, such that each segment represents an atomic execution. Such segmentation is highly preferred since some processes may block on some actions, such as the *wait* action. Therefore the processes need to be grouped into atomic sequences of execution, which are characterized by being cycle free (or at least containing zero time cycles only, depending upon the particular target language for implementation) and also by being free of any time consuming action, except at the very end of the sequence.

The segmentation process takes the sequential control flow graph, which is preferably (but not necessarily) a cyclic directed graph, and extracts all possible paths from the graph which start and finish with a wait node. Preferably, the entry and exit points to the whole graph are considered to be wait nodes too, as a simplification. Paths may fork and join, but may still be considered to lie within the same segment, as long as they are not interrupted by a time-consuming action such as a wait node.

Figure 9 shows an example of the segmentation of the client process for the sequential control flow graph. The segments are marked as numbers on the basic blocks. Note that each block may optionally be covered by multiple segments (such as the top node). Also, segments may optionally fork and join, as shown with regard to segment 4, which forks in this example.

According to optional but preferred embodiments of the present invention, the sequential control flow graph itself is more preferably optimized by eliminating exponential expansion of segments which contain nested *if-then-else* clauses. Such exponential expansion is eliminated by merging sequential control flow graph nodes after the segmentation phase, such that the merged nodes refer to the same source action and therefore represent the same action, and are also part of the same segment. One example of a suitable optimization process is performed by first computing a hash number for each node, based on the source reference and segment for that node; and then traversing each segment, replacing references to nodes with the first node in the node list that matches the hash value of the referenced nodes.

As an example, a simple program may have the following structure:

```
If B then {  
    A1  
} else {  
    A2  
};  
A3
```

The un-optimized control flow graph starts with a guard node, whose positive child would be a basic block node containing A1 followed by a basic block node containing A3. The negative child is a basic block containing A2 followed by a second instance of a basic block containing A3.

- 5 The optimization process stores the two basic blocks containing A3 with the same hash value. When re-constructing the control flow graph, the same instance of a basic block node will be used to represent A3 under both the positive and the negative branch. This reduces the size of the control flow graph.

- 10 As shown with regard to Figure 1, in the next phase, an interrelationship is created between the elaboration graph and the sequential control flow graph. Such an interrelationship considers the effects of each process instance on the elaborated structure, and thereby uncovers communication through global variables, side effects and race conditions.

- 15 In order to create such an interrelationship, preferably the control flow is first projected over the structure, by identifying all instances that have nested processes. These instances serve as the *me* (or the self reference) context for their processes. Processes which are present in the pre-run time structure are termed “root processes”, and are started through the *run()* method. Such processes typically have an infinite loop at their top level. Root processes may in turn call
20 methods or start new processes. The issues related to method calls are discussed in greater detail below.

Since each nested process in a list item breaks the list symmetry, the nested processes are optionally and more preferably classified into several

classes, each of which is then handled separately. A first class is the independent process, which only accesses data in the list object, but which does not access other list objects or the global context. This process does not actually break the symmetry, and one instance of such a process is sufficient to demonstrate the effects of all such processes on the elaboration graph. The next class is the neighbor dependent process, which uses its index to access other lists, and peeks into the predecessor and successor objects in the list. Multiple such processes need to be present in order to see the full effect on the elaboration graph.

For example, suppose that a list item can access fields within the object as well as fields within the predecessor object in the list. All list elements are symmetric, except for item 0 that has no predecessor. To fully evaluate the effect of such objects in a program, two cases must be considered: the case of item 0 and the case of an item whose index is greater than 0. In general neighbor dependent processes can be analyzed by considering more than one list item but less than the full list. Such processes break symmetry in some ways but may preserve it in other ways, such that a small set of representative processes may need to be considered in order to identify the full effect on the elaboration graph..

The third class is the arbitrarily dependent process, which is not restricted in access at all and therefore which always breaks symmetry. The effect of these processes cannot be determined by examining the behavior of a small set of such processes thereof, such that the presence of this type of process must be assumed to break symmetry and thus must be examined independently.

Next, the process of annotation optionally and more preferably handles aliasing, which is a problem introduced by dereferencing. If a variable, such as a parameter or field, is assigned the reference to a node in the elaboration graph, any operation performed on such a variable should preferably be applied to the
5 elaboration graph node. Therefore, the variable now functions as an alias to the node. However, if during the processing of the control structure such dereferencing is passed through a plurality of different assignments, some of which are conditional, the end result could be an indeterminate, non-defined aliasing relationship which cannot be determined during static analysis. In order
10 to solve this problem, preferably a list of possible aliases is maintained for each variable. If a decision to select a particular alias must be made during run-time, instantiated selector logic then preferably addresses the correct node.

As an example, consider the following simple program:

```
Struct f_type {  
15     Left:int;  
     Right:int;  
};  
  
Struct A {  
     F : f_type;  
20     B :bool;  
     Prog():int is {  
         If B then {  
             Result = f.left;
```

```
        } else {  
            result = f.right;  
        }  
    }  
5      }
```

The program *prog()* may return either *f.left* or *f.right*, depending on the value of *B*. The *result* variable is an alias for both. Some logic must be created in the target language, such that *result* will reflect the value of the correct field at runtime.

10 Such a solution is preferred, since it prevents unclear aliasing, but may have the disadvantage of requiring a large number of potential aliases to be stored.

As previously described, method calls may occur when the root process calls another method, for example. Preferably, each method call results in the actions of the called method being added to the sequential control flow graph, with formal parameters being replaced by actual parameters. Difficulties with this solution may arise from aliasing, as previously described, and polymorphism.

Polymorphic methods require the type to be resolved in order to determine which implementation to call, since in effect, the exact method is not determined until the particular implementation is selected. Such a selection of a particular implementation, and hence resolution of the polymorphism, may not be possible during static analysis for a number of reasons. For example, the static type may not be a leaf of the type tree, such that the actual type may be a subtype.

Alternatively, the static type may be an unresolved *when* sub-type, such that the determinant may change its value at run-time. In order to solve this problem, more preferably all possible cases are supported, with sufficient logic added to be able to select an appropriate implementation during run-time. The complexity
5 limitations are similar to that for the aliasing case, but these two effects may amplify each other, potentially inflating the size of the structure created in the target language.

Once these different types of ambiguities have been resolved, or at least sufficiently supported to enable resolution at run-time, the annotation process
10 preferably then continues with the tagging of all instances in the elaboration graph with labels that identify data access. A tag more preferably features the accessing instance, the accessing segment within the process instance and the access mode (read or write). The tagging process is optionally and more preferably performed as follows. First, the elaboration graph is traversed and
15 every process instance which is found preferably serves as a starting point. At each starting point, for every segment, all actions are preferably stepped through. Every data access expression of each action is mapped onto the elaboration graph, more preferably also according to the action context. Each target instance may then be identified, and tagged according to the semantics of the data access
20 expression. As the process of tagging is performed, a global data flow pattern emerges, as the information is used to compute a static schedule for the output model when possible, and optionally also in order to produce diagnostic messages for display to the user.

Figure 10 shows exemplary pseudocode for an exemplary annotation process with regard to tagging. This pseudocode emphasizes that the annotation starts with each node, and then penetrates more deeply into the structure of each node in a recursive manner, until each expression within each action is found and
5 tagged.

Turning again to Figure 1, the next stage is optionally and more preferably scheduling of execution for the output translated code in the target language. Such scheduling is highly preferred in order to maintain the scheduling rules of the original language, such as the *e* language, during execution in the target
10 language, which may either use different scheduling rules or which may actually execute in a truly concurrent environment. In such a concurrent environment, events do not occur according to a potentially artificial order or schedule, but instead occur as defined by the program. By contrast, if a target language executes according to different scheduling rules, events may occur according to
15 an arbitrary order.

According to an optional but preferred embodiment of the present invention, in order to prevent such arbitrary ordering, preferably scheduling considerations are compiled into the structure of the translated code in the form of guards. Thus, such considerations are more preferably required to be static in
20 nature, in order to enable these considerations to be immediately realized as guards.

The scheduling phase of the method of the present invention, as described with regard to Figure 1, therefore preferably attempts to find a static order of

priority for all segments, which is herein termed the schedule. Given an execution state, which is the list of values associated with all control variables, the schedule determines the order in which all exposed actions are performed. Exposed actions are herein defined as those actions with guards that are **true**.

- 5 A particularly preferred aspect of the scheduling phase is the elimination of races. A race condition is a very common and serious issue in circuit design. An example of a race condition is shown in Table 1 below by a Verilog program.

Table 1. Race Condition

```

10           module t(clk, a, b, c);
              input clk;
              output a;
              input b;
              input c;
              reg a;
15
              always @(posedge clk) a = b;
              always @(posedge clk) a = c;

              endmodule
20
```

In this example,

when the clock "clk" changes from "0" to "1", that is at the positive edge of clk, either $a = b$ may happen first or $a = c$ may happen first. That is, they may race each other. As a result, signal "a" may get the value of "b" or "c." In circuit design, this is an undesired effect because the final value of "a" is
5 nondeterministic.

A race condition is defined as a situation where reading a variable may result in one of several values nondeterministically. To avoid races, circuits are usually designed in some specific ways to prevent a race condition. Another approach is to detect during the design stage if the circuits may have races. Some
10 circuit simulation tools can detect a race condition when the simulator finds that, during the same simulation cycle, a signal (also referred to as a wire, register, or variable) is being written to two or more times, or being written to as well as being read from (see, for example, U.S., Patent Nos. 5,383,167; 5,818,263; 5,901,061; and 6,009,256).

15 If the design is race-free, the only constraint on the schedule is to place expression computation before the result of that computation is used. More preferably, a special case is the evaluation of temporal expressions which are required before any TCMs are evaluated. According to the preferred implementation of the present invention with the ϵ temporal language, events
20 may optionally be emitted by TCMs, and must then be permitted to influence the execution of the code within the same clock cycle. The chain of dependency may be handled for at least certain cases of nested dependencies for events emitted by TCMs, for example when an emit operation within one TCM causes another

TCM to execute. In order for such dependencies to be correctly handled, the execution of the emit operation is more preferably moved to the beginning of each cycle, a process which is termed "action retiming" herein. Although action retiming itself may cause a scheduling problem, this problem is resolvable within
5 each individual segment.

Optionally and more preferably, only cycle synchronous models are handled by such retiming processes and other scheduling rationalizations. A cycle synchronous model is a model in which the computational results of one cycle influence the next cycle, and is for example the model for the *e* language.
10 If cycle asynchronous events occur, the user is more preferably alerted in order to resolve these events manually.

Also optionally and more preferably, potential races are identified as part of the analysis of the code which is being translated. Generally, a race is the result of an ambiguous program. The ambiguity of the program causes the
15 runtime scheduler, which determines the timing and order of events, to take an arbitrary decision concerning the order of execution of these events. This order may not then replicated by the actual device under test (DUT). Thus, the user is preferably notified of all suspected races.

A more detailed description of races is that races are created by processes
20 communicating through shared variables without an adequate access protocol. At the annotation phase, each instance in the elaboration graph is tagged with all accesses. Each instance that has multiple access tags, in which at least one such access is a "write" access, may potentially be a suspected race target. However,

preferably access tags which belong to the same segment instance are not considered to be a race, as they belong to a single thread of control and therefore will execute sequentially.

More preferably, each suspected race target is examined with regard to the multiple actions which attempt to access that target. Most preferably, the hierarchy of guards for regulating each action is examined, such that the suspected race target is considered to be a race if there exists at least one condition under which these action are exposed. For example, optionally and most preferably, determining whether two boolean conditions are simultaneously true could be used to filter potential races. Also optionally and most preferably, temporal expressions are resolved in order to determine whether two segments can execute concurrently.

In order to at least reduce problems which may be caused by such races, according to preferred embodiments of the present invention, temporal expressions are given priority over TCMs for execution. Also, if read-write races are present, more preferably the write action is scheduled before the read action, in order for values to be computed before they are used. Other actions may also optionally be schedule in order to enable the result of any given action to be determined before this result is require by the next action.

These constraints on the schedule of execution are optionally and more preferably embodied in a scheduling constraint graph. The nodes of the graph are preferably each a segment instance, as such an instance is an atomic, zero-time, sequence of actions. Each constraint which requires the segment to execute

before another segment, because of any action of the former segment, is therefore preferably represented as a direct edge from the dependent segment to the independent segment.

Figure 11 shows a simple two process example as written in the *e* language. This example contains two objects of type *agent* nested under *sys*. These agents execute concurrently, attempting to write a value to a common variable, which is the field *w* under *sys*. The agents implement a mutual exclusion algorithm that prevents concurrent access to *w*, which may cause data corruption. The algorithm requires each agent to assert a request and deposit some type of identifier in a shared, unprotected variable (the field *K* under *sys*). If the value remains intact and the other agent deasserts it's request, then the shared variable *w* may be written to again.

The resultant elaboration graph is shown in Figure 12. The segmented control flow graph is presented in Figure 13, while the derived segmented scheduling constraint graph is shown in Figure 14.

The elaboration graph shows the two shared variables *w* and *K* under *sys*, as well as two instances of the agent object representing the two agents. Each agent has an ID field *id*, a request field *req* and a reference to the other agent *oa*.

The control flow graph shows that each agent can be in one of 3 execution states: segment 0 is the initial request. Segment 1 represents the wait cycle that may be entered if the other agent has access to the shared variable and segment 2 is the state where the agent has access to the shared variable.

The scheduling constraint graph shows the conflicts between the segments of the two agents. Each ellipse represent a segment of one of the agents, and each edge represent a dependency between segments. Co-dependent segments represent conflicts in the schedule. The dashed nodes in the example scheduling constraint graph form two cycles which are the manifestation of suspected write/write races with respect to **K** and **w**. The race with respect to **K** is a true race, while the race with respect to **w** is not a true race. In order to detect and remove the second cycle, filtering would need to be performed with regard to the temporal properties of the processes which are involved.

10 Once the scheduling constraint graph has been created, preferably the scheduling process creates a topological order over this graph. The process selects an unconstrained node to schedule, after which all of the constraints leading to that node are removed. This process is more preferably repeated until all of the nodes are scheduled. If the scheduling-constraint graph itself is cycle free, then the process may optionally be repeated without interruption.

15 However, if this graph has one or more cycles, then the process may halt if no additional independent nodes are available. The process more preferably then performs an arbitrary decision, such that some constraints are ignored. This arbitrary decision may result in poor timing of the output model, such that the user is most preferably notified when such a decision is taken. The presence of such a cycle in the scheduling constraint graph may be considered to imply corresponding timing ambiguity which had not been resolved as previously described. On the other hand, the presence of a cycle alone is not sufficient for

causing a race, since the corresponding constraints may have resulted from a false race.

Furthermore, once a node is scheduled, more preferably all the constraints which lead to this node are removed, which may break one or more cycles. The original process may then optionally continue. Most preferably, one or more
5 heuristic rules may be added, in order to minimize such cycle breaking, since cycle breaking may result in the introduction of a timing difference between the original and synthesized models.

As previously described, according to preferred embodiments of the
10 present invention, action retiming is performed in order to support emitted events within TCMs and therefore to rationalize behavior of the executed code. Since an event which is emitted by a TCM is expected to influence the flow of control within the same clock cycle, preferably retimed operations are extracted from the control flow at the annotation stage, more preferably with the guards which
15 govern their execution. The extraction process preferably determines whether the zero time path from the segment entry point to the retimed action contains actions that modify guard conditions, and the user is notified of the existence of any such actions which modify guard conditions.

Retimed actions are then more preferably placed first in the static
20 schedule, such that whenever an action would execute in a particular cycle, depending upon the values of its guards, that particular action is more preferably placed to execute first, before any segment is executed. The results of this execution are then ready to be used by the segments as necessary.

The previous description centered around an example of the method of the present invention with regard to any pair of an original language code and a target language, in which the code in the original language is translated to the target language. A particularly preferred example of an original language/target language pair is code written in the *e* language for the original language, and Verilog HDL as the target language.

While the invention has been described with respect to a limited number of embodiments, it will be appreciated that many variations, modifications and other applications of the invention may be made.